

1           **METHOD FOR TESTING OF A SOFTWARE EMULATOR WHILE**  
2           **EXECUTING THE SOFTWARE EMULATOR ON A TARGET MACHINE**  
3           **ARCHITECTURE**

4  
5           **BACKGROUND OF THE INVENTION**

6           Technical Field

7           The present invention relates to software emulators and, more particularly, relates  
8           to testing of software emulators.

9           Description of Related Art

10          An emulator is generally a device that is built to work like another. An emulator  
11          can be hardware, software, or a combination thereof. For example, an emulator may be  
12          software running on one machine that is designed to emulate another type of computer.  
13          Likewise, an emulator may be designed to execute software that was written to run on  
14          another computer. Typically, one computer emulates another computer (*i.e.*, a computer  
15          with a different instruction set architecture (ISA)). A computer's normal running mode,  
16          called native mode, is the operational state of the computer when the computer is  
17          executing programs from the computer's built-in instruction set (host machine  
18          "architecture"). In contrast, an emulation mode is the operational state of the computer  
19          while the computer is running a foreign program under emulation (using a "target  
20          machine architecture"). A foreign program is a program written in a different (foreign)  
21          ISA (*i.e.*, the target ISA) than the host machine's ISA (*i.e.*, the host ISA).

22          When an emulator is built, it is often necessary to test the emulator to determine  
23          the emulator's effectiveness. Typical and obvious methods of testing emulators have  
24          involved such strategies as: 1) forcing an emulated machine into a known state,  
25          executing a test program under emulation, and comparing a finished state vector in the  
26          emulator against a hand generated expected state vector; 2) forcing an emulated machine  
27          into a known state, executing a state program under emulation, and comparing the  
28          emulations finished state vector against the finished state vector of a previous version of  
29          the emulator (*i.e.*, "known good " finish state vector) – in a practice commonly known as  
30          "regression testing"; 3) forcing an emulated machine to a known state, executing a test  
31          program under emulation, and comparing the finished state vector against the finished  
32          state vector of the program as run on the target machine architecture; and, 4) forcing an  
33          emulated machine into a known state, executing a random sequence of machine

1 instructions under emulation, and comparing the emulator's finished state vector against  
2 the finished state vector of the same sequence as run on the target machine architecture.

3 The above approaches of testing emulators all require the test program to be run  
4 under emulation and the output of the emulation to be compared with something else – a  
5 hand generated prediction, output of another emulation run, etc. Figure 1 is a diagram  
6 generally illustrating the above approaches to testing emulators. Among the  
7 disadvantages of the above approaches is the likelihood of clerical error. For example,  
8 when comparing the output of one emulation run against another emulation run one  
9 output file may easily be confused with another. Likewise, hand-generated predictions of  
10 the output state are subject to transcription errors. Furthermore, in certain non-Unix  
11 operating systems such as Multi-Programming Executive ("MPE"), the state vector may  
12 contain important information, which, unfortunately, varies from one execution to  
13 another. Such variation makes simple-minded comparisons of state vectors difficult.  
14 MPE is a multi-user operating system (OS) developed by Hewlett-Packard Company in  
15 the 1970's. The current version of MPE is POSIX compliant and supports UNIX  
16 function calls.

17 What is needed, therefore, is a method of testing emulators that does not involve  
18 such complex and error-prone state vector setup and such complexity and susceptibility to  
19 errors in the checking of results.

## 20 **BRIEF SUMMARY OF THE INVENTION**

21 The objects, features and advantages of the present invention are readily apparent  
22 from the detailed description of the preferred embodiments set forth below, in  
23 conjunction with the accompanying Drawings in which:

24 The present invention is a method and system that overcomes the above problems  
25 by executing a test program on a target machine architecture, executing the test program  
26 under emulation on the same target machine architecture, and comparing the results of the  
27 two program executions to evaluate the emulator. One such situation where it is desirable  
28 to test an emulator is when a developer is developing a new ISA and a new Operating  
29 System (OS) for the new ISA. In current technology, programs written in a so-called  
30 high-level language (*e.g.*, C++, Java) can be "compiled" into multiple ISAs and can be  
31 confidently asserted to have substantially the same behavior. Therefore, when the  
32 developer is developing an OS for a new ISA, the developer can develop an emulator  
33 portion (written in a high-level language) of the new OS for the new ISA, compile the

1 emulator in the old ISA and test the emulator on the old ISA (using the old ISA as the  
2 host ISA and the target ISA).

3 For example, if the old ISA is "PA-RISC" and the new ISA is "XYZ", the  
4 developer could build an emulator for XYZ in a high-level language, compile the  
5 emulator in PA-RISC, and run the emulator on the PA-RISC itself (*i.e.*, PA-RISC is the  
6 host ISA). This procedure catches the majority of errors in the emulator and minimizes  
7 the number of errors upon re-compiling and executing the emulator in XYZ since the  
8 compiled program (the emulator) can be confidently asserted to have substantially the  
9 same behavior under PA-RISC as under XYZ. If the new ISA (*e.g.*, XYZ) is not built  
10 yet, then testing the emulator on the old ISA (*e.g.*, PA-RISC) can expedite the overall  
11 development process.

12 Another advantage of the present invention is that the present invention  
13 overcomes the disadvantages of the prior art. Another advantage of the present invention  
14 is that the present invention avoids susceptibility to errors present in prior methods of  
15 testing emulators. Another advantage of the present invention is that where an emulator  
16 is trying to emulate a target machine ISA, the method permits a test program to be run  
17 directly on the target machine ISA and also under emulation (on the same target machine  
18 ISA) within a single machine process (*i.e.*, the host machine ISA is the same as the target  
19 machine ISA). This advantage leads to another advantage of the present invention in that  
20 the present invention avoids the variation in state vectors that is a problem with the prior  
21 art methods of testing emulators.

22 These and other advantages of the present invention are achieved in a method for  
23 testing a software emulator while executing the software emulator on a target machine  
24 architecture, comprising the steps of executing a test program on a target machine  
25 architecture, executing an emulator on the target machine architecture, and the emulator  
26 executing the test program under emulation. Executing the test program produces a first  
27 output and the emulator executing the test program produces a second output.

28 These and other advantages of the present invention are also achieved in a  
29 computer readable medium comprising instructions for testing a software emulator while  
30 executing the software emulator on a target machine architecture, by executing a test  
31 program on a target machine architecture, whereby the test program produces a first  
32 output, executing an emulator on the target machine architecture, and the emulator  
33 executing the test program under emulation, whereby the test program produces a second  
34 output.

1        These and other advantages of the present invention are also achieved in a  
2 computer readable medium containing a program that includes instructions for testing a  
3 software emulator while executing the software emulator on a target machine  
4 architecture, by executing a test program on a target machine architecture, whereby the  
5 test program produces a first output and executing an emulator on target machine  
6 architecture. The emulator calls the test program and executes the test program under  
7 emulation, whereby the test program produces a second output.

## 8    **BRIEF DESCRIPTION OF THE DRAWINGS**

9        The invention will be described, by way of example, in the description of  
10 exemplary embodiments, with particular reference to the accompanying drawings, in  
11 which like reference numbers refer to like elements, and in which:

12        Figure 1 is a prior art diagram illustrating conventional emulation testing methods.

13        Figure 2 is a flowchart illustrating an embodiment of a method for testing  
14 emulators while executing the emulator on a target machine architecture.

15        Figure 3 is a block/flow diagram illustrating an embodiment of a system method  
16 for testing of a software emulator while executing the software on a target machine  
17 architecture.

18        Figure 4 is a flow chart illustrating steps of emulator execution testing in detail.

19        Figure 5 is a block diagram of a host machine on which the system and method for  
20 testing a software emulator may be executed.

## 21    **DETAILED DESCRIPTION OF THE INVENTION**

22        In a preferred embodiment of the present invention, a single test program  
23 produces all the output necessary for testing an emulator and therefore various clerical  
24 errors present in the prior art are rendered irrelevant. In a preferred scenario of usage, a  
25 test program (possibly a random sequence of machine instructions) is embodied in a  
26 subprogram or subroutine. The main program performs a subprogram or subroutine call  
27 to execute the test program directly on a target machine architecture. When the test  
28 program is finished executing, the main program then causes the test program to execute  
29 under emulation. Since both the emulated execution and a direct execution of the test  
30 program execute within the same process, certain kinds of process-to-process variations  
31 do not occur, and the comparison described above is reduced to a comparison between the  
32 "first half" and the "second half" of the single program's output. Executing the test

1 program directly and under emulation within the same process simplifies the clerical  
2 aspect of testing by removing several sources of error.

3 A preferred embodiment of the present invention includes components, such as an  
4 emulator, which executes on a target machine architecture, in the form of a callable  
5 library, for example. The emulator preferably includes a callable entry point having  
6 semantics "begin emulation at the return point of this subroutine." Such an entry point,  
7 when invoked, preferably copies the host machine's state vector into the emulated  
8 machine's state vector, and emulates instructions starting at a point where the invoking  
9 program (the main program) would have returned (*i.e.*, the next instruction in the main  
10 program). Another component of the preferred embodiment is a subprogram to be tested  
11 (a "test program"), which will be executed directly upon the target architecture and also  
12 under emulation (using the same target machine architecture). The effects of the  
13 subprogram are preferably visible in a standard output or default print output stream.  
14 Another component of the preferred embodiment is preferably a main program that may  
15 be coded as follows:

```
16 1001 CALL sut
17 1002 CALL es
18 1003 CALL sut
19 1004 MOV %arg0,#0
20 1005 CALL _exit,
```

21 where "CALL" does a subroutine call, "MOV" causes a register to be set to a certain  
22 value (the "arg0" register set to zero in this case), "sut" is an address of an entry point of  
23 the subroutine under test (the test program), "es" is an address of an emulation  
24 subsystem's (the emulator's) entry point, and "\_exit" is an address of a supervisor call to  
25 terminate execution.

26 In the above situation, the host machine executes the first "CALL sut", then  
27 executes the instructions in the test program. The last instruction of the test program  
28 executed is typically a "RETURN" type of instruction, which basically means "resume  
29 execution at the point (in the main program) after the CALL instruction that started the  
30 test program."

31 Then the host machine executes the "CALL es" instruction, and then starts  
32 executing the emulator. The emulator determines that the emulator was called from the  
33 address "1002". The emulator then determines that the next instruction to emulate is at  
34 1003, and the emulator's very first "next instruction" is the 2nd "CALL sut" instruction at  
35 "1003" above. The subsequent instruction is the first instruction of the test program.

1 Again, the very first "next instruction" is the "CALL sut" (if there is additional  
2 setup, it would include that setup as well) – the next instruction after that would be the  
3 first instruction of the subroutine under test.

4 The preferred embodiment may also include a further component such as an  
5 automated comparison subroutine that compares the first output produced by the first  
6 execution of the test program to a second output produced by the second execution of the  
7 test program (the execution of the test program under emulation).

8 Figure 2 is a flowchart illustrating a method 10 for testing the software emulator  
9 while executing the software emulator on a target machine architecture. The method 10  
10 preferably comprises: calling a test program 12; executing a test program on a target  
11 machine architecture 14; calling an emulator 16; executing the emulator on the target  
12 machine architecture 18; the emulator calling the test program 20; the emulator executing  
13 the test program under emulation 22; comparing a first output produced by the executed  
14 test program to a second output produced by the test program executed under emulation  
15 24; and, determining if the second output is within a certain margin of variation from the  
16 first output 26.

17 Calling a test program 12 preferably comprises the main program executing a  
18 native call of the test program. Calling a test program 12, therefore, triggers the  
19 executing of the test program 14, in a native mode, on the target machine architecture. In  
20 this situation, the target machine architecture is a basic set of instructions for machine  
21 language with which a host machine running a main program is coded. Executing the test  
22 program of a target machine architecture 14 preferably comprises the test program  
23 executing the test program instructions until the test program reaches an end of program.  
24 The executed test program instructions preferably produce at least one output ("a first  
25 output"). The test program instructions may produce a plurality of outputs when executed  
26 (i.e., the first output is the plurality of outputs). When the test program reaches a  
27 "RETURN", a native return is executed returning control to the main program.

28 Calling an emulator 16 preferably comprises the main program calling the  
29 emulator. The calling of the emulator 16 is preferably executed as a native call. Calling  
30 the emulator 16 triggers executing the emulator on the target machine architecture 18.  
31 Executing the emulator on the target machine architecture 18 preferably comprises the  
32 emulator setting up a state vector and then executing the next instruction(s) of the main  
33 program. The emulator, therefore, performs an emulated return to the main program to

1 execute the next instruction of the main program. The next instruction on the main  
2 program may include setup instructions such as printing a "begin emulation mode"  
3 message. The emulator executes such setup instructions and then executes the next main  
4 program instruction, which preferably calls the test program. Therefore, the emulator  
5 calling the test program 20 preferably comprises the emulator executing a main program  
6 instruction to affect an emulated call of the test program. As described above, the  
7 instruction set architecture of the emulator is the same as the host machine instruction set  
8 architecture. The emulator calling the test program 20 triggers the emulator executing the  
9 test program under emulation 22.

10 The emulator executing the program under emulation 22 preferably comprises the  
11 emulator executing each instruction of the test program in sequence until a "RETURN" or  
12 other end of program is reached. As noted previously, the test program preferably  
13 includes an instruction that produces at least one output. When the emulator executes the  
14 test program, the execution under emulation preferably produces a "second output". The  
15 second output may include a plurality of outputs, as above. When the test program  
16 reaches the "RETURN", an emulated return is preferably performed to return control to  
17 the main program. The emulator then executes the next instruction of the main program

18 Comparing the first output to the second output 24 preferably comprises an  
19 automatic or automated comparison of the first output produced during the first execution  
20 of the test program to the second output produced by the execution of the test program  
21 under emulation. The automated comparison may be encoded in the main program or  
22 separately as a subroutine or program. Determining if the second output is within a  
23 certain margin of variation of the first output 26 preferably comprises the automated  
24 subroutine program calculating the difference between the first output and the second  
25 output and determining whether the difference is within a pre-set margin of variation. If  
26 the difference is within the pre-set margin of variation, then the emulator is performing  
27 satisfactorily. The method 10 of testing may be repeated as necessary to produce a  
28 sufficient sample of results. Likewise, the method 10 may be repeated after  
29 changes/corrections are made to the emulator and the emulator is re-compiled in the  
30 target ISA (the host ISA).

31 Figure 3 is a block-flow diagram illustrating an emulator 32, main program 34,  
32 and a test program 36. The solid arrows in Figure 3 represent native transfers of control  
33 while the dashed arrows represent emulated transfers of control. The main program 34  
34 preferably comprises a series of instructions as described above. The instructions of the

main program preferably comprise a test program call 38 that executes a native call of the test program 36. The test program 36 preferably comprises a series of instructions, as described above. The test program 36 instructions preferably produce an output(s). The test program 36 preferably executes a native return to the main program 34 when the test program reaches a "RETURN" or other end of program information. The main program 34 preferably also includes a call emulator instruction 40 that preferably executes a native call of the emulator 32. The emulator 32 preferably includes a state vector 42 that is loaded with the state of a central processing unit ("CPU") on which the main program 34, test program 36, and emulator 32 execute. The main program 34 includes a portion, shown with cross-hatching in Figure 3, that represents an emulated execution. The emulated portion includes a second test program call instruction 44 that is executed by the emulator 32 to perform an emulated call of the test program 36. After the test program 36 executes under emulation and performs an emulated return to the main program 34, the emulator 32 executes a RETURN instruction. As shown in Figure 3, the emulator 32 does not have a separate "memory" portion of its state vector 42, rather, the emulator 32 simply uses a native memory of the CPU. A comparison subroutine or program is not shown in Figure 3.

Figure 4 is a flow diagram illustrating steps 18 through 22 of the method 10 in greater detail. A method 50 of emulator execution and testing shown in Figure 4 comprises: calculating and saving the target machine's state at the point of return 52; allocating additional resources to execute the emulator 54; reading the next instruction according to the state vector 56; determining whether the next instruction requires a supervisor call 58; if a supervisor call is not required, updating the state vector according to the next instruction 60; if a supervisor call is required, executing the supervisor call 62 and updating the state vector according to the result of a supervisor call 64. As shown in Figure 4, the method 50 will repeat until a supervisor call that does not return (e.g., an "\_exit" supervisor call) is executed. When an "\_exit" supervisor call is executed, the process (including the emulator) terminates as a normal native program would.

Calculating and saving the target machine's state at the point of return 52 preferably comprises the emulator 32 determining the instructions of the main program following the point at which the emulator 32 was called by the main program and loading these instructions into the state vector in the emulator 32. When the emulator 32 is called, the call typically causes certain parts of the host machine CPU's state to be saved in a memory "stack". A subroutine needs to know where this information is saved (i.e., to the



1 extent necessary to execute the “return from subroutine” instruction). The emulator 32,  
2 itself coded as a subroutine, must thus know where this information is stored. The  
3 emulator 32 is preferably programmed to assign the address of the first instruction in the  
4 main program that will be executed when the emulator 32 returns to the main program to  
5 a variable (*e.g.*, the “return address”). The emulator 32 can then read the contents of the  
6 instruction memory using the return address to load the first instruction into a variable in  
7 order to emulate the first instruction.

8 As described above, when a subroutine (of which the emulator is one) is called,  
9 the return address is typically written into (or “pushed onto”) the stack. When the  
10 subroutine returns to its caller, the return address is read from (or “popped off”) the stack.  
11 The subroutine may also allocate local variables (temporary storage) on the stack. When  
12 the subroutine returns to its caller, the local variables are deleted from the stack. There  
13 may also be a “heap” from which memory may be allocated in a less transient manner. A  
14 subroutine may, for example, allocate a block of memory from the heap, write a pointer to  
15 that block into a global variable, and return. Whereas the subroutine’s local variables all  
16 are deleted when the subroutine returns, the memory allocated from the heap remains.

17 The emulator 32, being itself a subroutine that may also call other subroutines,  
18 typically uses the stack and the heap in the manner described above. That the emulator  
19 32 preferably never actually returns to the main program does not detract from the fact  
20 that in an emulated return to the main program, the resources allocated by the emulator 32  
21 must be, in the emulated machine, returned for use by the emulated program.

22 A program being emulated (*e.g.*, the test program) will typically also make use of  
23 the stack and the heap. Therefore, unless precautions are taken, the emulator’s use of the  
24 stack and the heap may conflict with the test program’s use of the stack or the heap.  
25 Therefore, allocating additional resources to execute the emulator 54 preferably  
26 comprises the emulator 32 allocating a separate stack and/or heap for its own use. Certain  
27 architectures may also require additional resources.

28 Reading the next instruction according to the state vector 56 preferably comprises  
29 the emulator reading out the next stored instruction in the state vector. When the  
30 emulator is first called, the next instruction is the next instruction of the main program. If  
31 the next instruction of the main program calls the test program, the subsequent instruction  
32 will be the first instruction of the test program.

33 In the determining step 58, the emulator 32 determines if the next instruction of  
34 the main program requires a supervisor call. A supervisor call is a mechanism used by an

1 application program to request service from the operating system. System calls often use  
2 a special machine code instruction that causes the host machine CPU to change mode  
3 (e.g., to a supervisor mode). The supervisor mode allows the host machine operating  
4 system to perform restricted actions such as accessing hardware devices or a memory  
5 management unit. If an instruction requiring a supervisor call is encountered, *i.e.*, if the  
6 emulator 32 cannot execute the instruction because the emulator is not designed to  
7 emulate the restricted actions, then the emulator 32 will perform a supervisor call 62 on  
8 behalf of the emulated program.

9 If the next instruction does not require a supervisor call, updating the state vector  
10 according to the next instruction 60 preferably comprises the emulator 32 executing the  
11 present instruction and returning to read the next instruction after the present instruction  
12 (*i.e.*, the emulator 32 repeats step 56 for the next instruction after the present instruction  
13 and so on until an `_exit` supervisor call is reached). If the next instruction is the first  
14 instruction of the test program (*i.e.*, the test program is called by the previous instruction)  
15 the emulator 32 proceeds to execute the instructions of the test program until a  
16 "RETURN" instruction in the test program is reached. The "RETURN" instruction  
17 returns control to the main program, causing the emulator 32 to resume execution at an  
18 instruction after the call test program instruction. The emulator 32 proceeds to execute  
19 remaining instructions of the main program (including, for example, a comparison  
20 subroutine) until it reaches an `_exit` supervisor call in the main program. When the  
21 emulator 32 reaches the `_exit` supervisor call, the emulator executes the supervisor call  
22 62 necessary to exit the main program and the main program is exited.

23 The objects of this invention may be accomplished utilizing components known in  
24 the art. Furthermore, the objects of the present invention may be accomplished by the  
25 operation of computerized system components implementable in hardware or software or  
26 in combination. Accordingly, a computer readable medium storing a program or  
27 containing instructions for realizing the objects of the present invention may be produced  
28 and is disclosed. Likewise, a processor with a memory containing instructions for  
29 realizing the objects of the present invention may be produced and is disclosed.  
30 Consequently, Figure 5 illustrates host machine 90, as described above, comprising a  
31 CPU or processor 92, a memory 94, a secondary storage 96, and an output device 98. As  
32 stated, the memory 94 (*e.g.*, RAM) and/or secondary storage 96 (*e.g.*, a hard-drive, CD-  
33 ROM, carrier wave) may be computer readable mediums and may store programs or  
34 contain instructions, executed by the CPU 92, for performing the above-described

1 methods and functions. For example, the secondary storage 96 or the memory 94 may  
2 store the test program 36, main program 34 and emulator 32 described above. The output  
3 device 98 (*e.g.*, a display, printer, speaker) may display or otherwise output the output of  
4 the test program 36.

5 The foregoing description of the present invention provides illustration and  
6 description, but is not intended to be exhaustive or to limit the invention to the precise  
7 one disclosed. Modifications and variations are possible consistent with the above  
8 teachings or may be acquired from practice of the invention. Thus, it is noted that the  
9 scope of the invention is defined by the claims and their equivalents.

09909057-071901  
T06T20-25060660